

1 METHOD AND SYSTEM FOR USING A BAYESIAN BELIEF NETWORK TO
2 ENSURE DATA INTEGRITY
3

4 BACKGROUND OF THE INVENTION

5 Field of the Invention

6 The present invention relates to a system and method for measuring the
7 financial risks associated with trading portfolios. Moreover, the present invention
8 relates to a system and method for assuring the integrity of data used to evaluate
9 financial risks and/or exposures.
10

11 Description of the Related Art

12 As companies and financial institutions grow more dependent on the global
13 economy, the volatility of currency exchange rates, interest rates, and market
14 fluctuations creates significant risks. Failure to properly quantify and manage risk
15 can result in disasters such as the failure of Barings ING. To help manage risks,
16 companies can trade derivative instruments to selectively transfer risk to other parties
17 in exchange for sufficient consideration.

18 A derivative is a security that derives its value from another underlying
19 security. For example, Alan loans Bob \$100 dollars on a floating interest rate. The
20 rate is currently at 7%. Bob calls his bank and says, "I am afraid that interest rates
21 will rise. Let us say I pay you 7% and you pay my loan to Alan at the current floating
22 rate." If rates go down, the bank makes the money on the spread (the difference
23 between the 7% float rate and the new lower rate) and Bob is borrowing at a higher
24 rate. If rates rise however, then the bank loses money and Bob is borrowing at a
25 lower rate. Banks usually charge a risk/service fee, in addition, to compensate it for
26 the additional risk.

27 Derivatives also serve as risk-shifting devices. Initially, they were used to
28 reduce exposure to changes in independent factors such as foreign exchange rates and

1 interest rates. More recently, derivatives have been used to segregate categories of
2 investment risk that may appeal to different investment strategies used by mutual fund
3 managers, corporate treasurers or pension fund administrators. These investment
4 managers may decide that it is more beneficial to assume a specific risk characteristic
5 of a security.

6 Derivative markets play an increasingly important role in contemporary
7 financial markets, primarily through risk management. Derivative securities provide
8 a mechanism through which investors, corporations, and countries can effectively
9 hedge themselves against financial risks. Hedging financial risks is similar to
10 purchasing insurance; hedging provides insurance against the adverse effect of
11 variables over which businesses or countries have no control.

12 Many times, entities such as corporations enter into transactions that are based
13 on a floating rate, interest, or currency. In order to hedge the volatility of these
14 securities, the entity will enter into another deal with a financial institution that will
15 take the risk from them, at a cost, by providing a fixed rate. Both the interest rate and
16 foreign exchange rate derivatives lock in a fixed rate/price for the particular
17 transaction one holds.

18 Consider another example. If ABC, an American company, expects payment
19 for a shipment of goods in British Pound Sterling, it may enter into a derivative
20 contract with Bank A to reduce the risk that the exchange rate with the U.S. Dollar
21 will be more unfavorable at the time the bill is due and paid. Under the derivative
22 instrument, Bank A is obligated to pay ABC the amount due at the exchange rate in
23 effect when the derivative contract was executed. By using a derivative product,
24 ABC has shifted the risk of exchange rate movement to Bank A.

25 The financial markets increasingly have become subject to greater "swings" in
26 interest rate movements than in past decades. As a result, financial derivatives have
27 also appealed to corporate treasurers who wish to take advantage of favorable interest
28 rates in the management of corporate debt without the expense of issuing new debt

1 securities. For example, if a corporation has issued long term debt with an interest
2 rate of 7 percent and current interest rates are 5 percent, the corporate treasurer may
3 choose to exchange (i.e., swap) interest rate payments on the long term debt for a
4 floating interest rate, without disturbing the underlying principal amount of the debt
5 itself.

6 In order to manage risk, financial institutions have implemented quantitative
7 applications to measure the financial risks of trades. Calculating the risks associated
8 with complex derivative contracts can be very difficult, requiring estimates of interest
9 rates, exchange rates, and market prices at the maturity date, which may be twenty to
10 thirty years in the future. To make estimates of risk, various statistical and
11 probabilistic techniques are used. These risk assessment systems—called Pre-
12 Settlement Exposure (PSE) Servers—are commonly known in the art.

13 PSE Servers often simulate market conditions over the life of the derivative
14 contracts to determine the exposure profile representing the worst case scenario
15 within a two standard deviation confidence interval, or approximately 97.7%
16 confidence. Thus, the PSE Server outputs an estimate of the maximum loss that the
17 financial institution will sustain with a 97.7% chance of being correct. This exposure
18 profile is calculated to give current estimates of future liabilities. As market
19 conditions fluctuate from day to day or intra-day, the calculated exposure profile
20 changes; however, these changes are not always due to market fluctuations, they are
21 sometimes due to errors in the input data.

22 23 BRIEF SUMMARY OF THE INVENTION

24 In the past, input data errors have been manually detected by credit analysts;
25 however, because the quantity of input data is so large, it is impractical for credit
26 analysts to detect and correct all of the errors. Credit analysts are most likely to detect
27 errors in the input data that cause a significant change in the exposure profile.

1 The Pre-Settlement Exposure (PSE) Server takes as input large amounts of
2 transactions and market data and in turn produces a significant amount of data and the
3 question is: Are the changes in the outputs due to a) the normal operation of the
4 system involving statistical simulation, b) expected market fluctuations, c) business
5 operations, d) system fault, or e) bad data. Thus, the accuracy of exposure reporting
6 by the PSE Server depends on the precision of its analytics and the quality of the data.
7 However, the data quality is not guaranteed and is difficult to test for every
8 permutation. Yet experience indicates that systematic validation must be
9 implemented because the possibility of artificially understating or overstating
10 exposure can adversely impact the business.

11 Nevertheless, the large volume and complex nature of derivatives transactions
12 and market data as well as the time constraints required to meet daily reporting
13 deadlines virtually preclude manual inspections of the data. It is possible in principle
14 to check every contract, every yield curve, or every exchange rate for they are inputs
15 to the PSE Server. However, because of reporting deadlines and the pace of business,
16 in practice this is not feasible on an intra-day or day-to-day basis. Thus, it is
17 convenient to treat the Server as a black box in terms of understanding all the causes
18 and effects that go into its operation.

19 The price to be paid for the black box perspective is that changes in
20 counterparty exposure sometimes seem unexplainable, even mysterious. A
21 counterparty is herein referred to a customer with whom there is some credit risk
22 (e.g., the risk that the customer may not pay what is owed at some future date.) Even
23 with a robot for automated verification analysis of the black-box Server to assist,
24 there remains a notable number of anomalous exposure shifts which escape the drill-
25 through analysis and consequently go "unexplained." Yet there must be a logical
26 explanation, only there are rarely human resources to regularly pursue it except when
27 a crisis arises or a problem becomes so intolerable the "experts" (such as credit

1 administrators, systems programmers, etc.) must be called in to sift through all the
2 data. The goal is to find a credible explanation from a) through e) above.

3 Nevertheless, this goal is not a simple task and in any event an enormous
4 distraction and drain of resources that could otherwise be focused on more important
5 business. If this process can be automated, at least for initial screening purposes,
6 there is considerable opportunity for savings of staff time and improving productivity
7 and end-to-end quality.

8 Hence, the preferred embodiments of the present invention provide a system
9 and method for a customizable Bayesian belief network to diagnose or explain
10 changes in the exposure profile of a risk assessment system, such as the Pre-
11 Settlement Exposure (PSE) Server, by performing induction, or backward reasoning,
12 to determine the most likely cause of a particular effect.

13 The preferred embodiments of the present invention further provide a method
14 and system for identifying plausible sources of error in data used as input to financial
15 risk assessment systems.

16 The preferred embodiments of the present invention further provide a method
17 and system for implementing a Bayesian belief network as a normative diagnostic
18 tool to model the relationship between and among inputs/outputs of the risk
19 assessment system and other external factors.

20 The preferred embodiments of the present invention also provide a system and
21 method for a Deep Informative Virtual Assistant (DIVA), which includes an
22 automated normative, diagnostic tool designed to use a Bayesian belief network (also
23 known as "Bayesian network") to "explain" changes in the exposure profile of a risk
24 assessment system such as a PSE Server.

25 The preferred embodiments of the present invention further provide a system
26 and method for a DIVA that provides sensitivity analysis and explanation context by
27 indicating the relative importance of an explanation in relation to an alternative
28 explanation.

1 The preferred embodiments of the present invention further provide a system
2 and method for a DIVA that is fast in mining data and interacting with the expert.
3 Thus, there is no perceptible degradation in performance of the normal processing
4 times on the PSE Server, and the interactive response time is short per query per
5 counterparty.

6 The preferred embodiments of the present invention also provide a system and
7 method for a DIVA that self diagnoses the explanation in terms of conflicts and
8 contradictions.

9 The preferred embodiments of the present invention further provide a system
10 and method for a DIVA that includes program modules, knowledge bases, statistical
11 history, and constraints for performing deeper analysis of data. Its knowledge bases
12 also contain detailed graphical information about causes and effects which allows the
13 system to make plausible inferences about systems and processes outside the PSE
14 Server “over the horizon” in both space in time.

15 The preferred embodiments of the present invention also provide a system and
16 method for a DIVA that supports the volume, complexity, and multifaceted nature of
17 the financial derivatives information processed by the PSE Server and performs
18 logical, systematic analysis of data integrity on such information.

19 The preferred embodiments of the present invention further provide a system
20 and method for a DIVA that is consistent for each counterparty and scalable at least
21 with respect to the number of deals and amount of market data.

22 The preferred embodiments of the present invention also provide a system and
23 method for a DIVA that is capable of making inferences “over the horizon” in both
24 space and time to point to potential sources of problems outside the PSE Server. The
25 DIVA is also capable of making predictions about future plausible outcomes given a
26 state of knowledge.

27 The preferred embodiments of the present invention also provide a system and
28 method for a DIVA that is designed in such a way that the contents and design of the

1 knowledge base is independent of the inference engine; thus, DIVA can be modular
2 for flexible modification.

3 The preferred embodiments of the present invention further provide a system
4 and method for a DIVA having at least three operational modes: (a) pre-release, (b)
5 post-release or follow up, and (c) preventative maintenance. Pre-release includes a
6 mode after a feed has arrived but before the hold-release decision is made by the
7 credit analyst. Post-release includes a mode after the hold-release decision is made
8 when credit analysts are expected to further investigate a run. Finally, preventative
9 maintenance includes a mode which is invoked periodically to scrub the system's
10 data, looking for potential problems ignored or suppressed during pre-release or post-
11 release modes. Each of these modes may also employ different standards of evidence
12 used to filter the analysis.

13 The preferred embodiments of the present invention also provide a system and
14 method for a DIVA that is configurable to explain production or quality assurance
15 (QA) environments. In fact, since normally find (or expect to find) many more
16 problems in QA, the system may have more utility here.

17 Additional aspects and novel features of the invention will be set forth in part
18 in the description that follows, and in part will become more apparent to those skilled
19 in the art upon examination of the disclosure.

20 BRIEF DESCRIPTION OF THE DRAWINGS

21 The preferred embodiments are illustrated by way of example and not
22 limitation in the following figures, in which:

23 Fig. 1A depicts the Pre-Settlement Exposure (PSE) server as a black box with
24 inputting causes and outputting effects in accordance to an embodiment of the present
25 invention.
26

27 Fig. 1B depicts the PSE server as a black box having each outputting effect
28 linked to an inputting cause in accordance to an embodiment of the present invention.

Fig. 2 depicts a Bayesian belief network in accordance to an embodiment of the present invention.

Fig. 3 depicts an architecture for a Deep Information Virtual Assistant (DIVA) in accordance to an embodiment of the present invention.

Fig. 4 depicts the name space relationships in a Bayesian belief network as implemented by a third-party software in accordance to an embodiment of the present invention.

Fig. 5 depicts a general architecture for a DIVA in accordance to an embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS OF THE INVENTION

Referring now in detail to an embodiment of the present invention, the system and method for a Deep Informative Virtual Assistant (DIVA), which make use of customized Bayesian belief networks (also known as “Bayesian networks”) to perform logical, systematic analysis of data integrity for risk assessment systems, such as Pre-Settlement Exposure (PSE) Servers, to ensure accurate evaluation of financial risks or exposures based on such information.

As is commonly known in the art, a Bayesian network works on the principle of Bayes’ theorem, named after Thomas Bayes, an 18th century Presbyterian minister and member of the British Royal Society. It is a knowledge base which is both structural and quantitative. The structural part is represented by a graph or network of nodes that describe the conditional relationships among variables in the problem domain. The quantitative part is represented by conditional probabilities that can be interpreted as the strengths of connections in the network.

According to an embodiment of the present invention, the PSE Server is a complex system with thousands of functions points. It takes as input financial information that fluctuates according to world market conditions. It also uses a

1 statistical process, such as the Monte Carlo simulation, to estimate realistic market
2 scenarios in the future. The Monte Carlo method provides approximate solutions to a
3 variety of mathematical problems relating to risk estimation and exposure-profile
4 generation by performing statistical sampling experiments. The method can be
5 applied to problems with no probabilistic content as well as those with inherent
6 probabilistic structure.

7 Because the PSE Server receives, analyzes, and generates large volumes of
8 transactions and market data, it is practically impossible to check each and every
9 datum. Thus, according to an embodiment of the present invention, it is convenient to
10 treat the PSE Server as a black box in terms of understanding all the causes and
11 effects that go into its operation. Figs. 1A and 1B depict the PSE server as a black
12 box with outputting effects associated with corresponding inputting causes.

13 Consequently, the essential problem is one of finding a needle in the haystack
14 because most of the data received and generated by a PSE server is correct.
15 Moreover, when there are significant changes in the data which usually cause
16 significant changes in the exposure profile, these situations are generally obvious.
17 Thus, it's the subtler, deeper problems that need to be discovered and corrected. By
18 logical analysis, prior experience, and common sense, the DIVA according to one
19 embodiment of the present invention is capable of finding the needle in the haystack.
20 In other words, DIVA is capable of reliably relating specific causes to specific effects
21 in the PSE server that saves staff time and resources.

22 While a risk assessment system, such as the PSE Server, can be treated as a
23 black box according to the preferred embodiments of the present invention, it is
24 expected to exhibit certain patterns of behavior according, informally, to the 80-20
25 rule. Namely, most problems are caused by a relatively few situations. For the
26 reasons given above, the connection between cause and effect is not typically
27 deterministic but probabilistic. As is known in the art, with a deterministic model,
28 specific outcomes of an experiment can be accurately predicted; whereas, with a

1 probabilistic model, relative frequencies for various possible outcomes of the
2 experiment can be predicted but not without uncertainty.

3 The connections between causes and effects and their strength in terms of
4 probability, as determined by DIVA, are represented in a knowledge base called a
5 Bayesian belief network. According to one embodiment of the present invention, the
6 belief network includes a graph capable of representing cause-effect relationships and
7 decision analysis that allows an inference engine to reason inductively from effects to
8 causes. Hence, as an automated but "supervised" assistant based on a belief network,
9 DIVA is intended to support rather than replace the credit analyst.

10 In one embodiment of the present invention, a third party software package,
11 such as the Hugin™ software, may be used to provide a Graphical User Interface
12 (GUI) shell for developing belief networks and an Application Program Interface
13 (API) for embedded applications. This software is herein referred to as the API
14 software. This software does not generate artificial intelligence. Rather, its main job
15 is to calculate the joint probability table,

$$P(X_1, X_2, \dots, X_N)$$

16
17 which would require $O(2^N)$ complexity for variables with just two states. For any
18 realistic N, say $N \approx 100$, a direct implementation of this table exceeds the capacity of
19 computers in service today and on the horizon for the foreseeable future. Yet without
20 actually generating the full joint probability table, the belief network, implemented by
21 API according to an embodiment of the present invention, can normally manage this
22 problem efficiently using various mathematical methods and system techniques via
23 software implementation that make use of more reasonable space and time.

24 According to an embodiment of the present invention, DIVA provides
25 infrastructure supports, both conceptually and in software, which interfaces with the
26 belief network. To that extent, at least one "expert" is employed to specify the
27 knowledge base in the form of a belief network for DIVA, wherein the belief network
28 represents a closed world of knowledge. Automated learning techniques may also be

1 applied to automatically generate the knowledge base. DIVA is then used to interpret
2 the results from the belief network. Indeed, one of the problems faced and resolved
3 by DIVA is the question of what constitutes "evidence" that a change of significance
4 has been observed when, as mentioned earlier, most of the time the data is correct.
5 The fact that there may be a problem embedded within a much larger collection of
6 correct data is the haystack. However, this fact can be seen as an advantage.
7 According to an embodiment of the present invention, the initial probabilities of the
8 Bayesian belief network can be set to reflect this experience, as explained in detail
9 later.

10 According to an embodiment of the present invention, DIVA's job includes
11 extracting the needle, i.e., identifying the source that plausibly accounts for the
12 problem. According to the present invention, plausibility refers to the existence of a
13 residue of uncertainty with any given assessment. Even if DIVA cannot find a
14 problem, it can rule out sources that are not likely causing the problem, which
15 remains useful to know in assessing the cause of an effect.

16 Because the belief network represents a closed world of knowledge, there
17 arises the possibility of logical contradictions. According to an embodiment of the
18 present invention, the idea of the closed-world representation of the belief network is
19 that DIVA conforms to Gödel's incompleteness theorem. As is known in the art, the
20 Gödel's incompleteness theorem limits what a system can do. That is, within any
21 logical system, there exists propositions that can neither be proved nor disproved.
22 Hence, any attempt to prove or disprove such statements by the defined rules within
23 the boundary of the system may result in contradiction. Accordingly, for DIVA to
24 conform to Gödel's incompleteness theorem, it would mean for all practical purposes
25 that DIVA either a) finds the cause for an effect with certainty, i.e., probability 1, or
26 b) contradicts itself.

27 A contradiction does not indicate that DIVA fails to function properly.
28 Indeed, if a Bayesian belief network produces a contradiction, DIVA indicates that it

1 is in this state and can thus inform the credit analyst. A contradiction can mean (a)
2 the inference engine that drives the belief network such as the API software or DIVA
3 has a bug that needs to be fixed; (b) more likely that the belief network is either truly
4 contradictory, in which case there is a bug in its design that needs to be fixed; or (c)
5 more likely that the network is incomplete. If the network is incomplete, that, too, is
6 useful to know because it provides information needed to bring the hypothesis space
7 of the knowledge base more in line with actual experience.

8 According to an embodiment of the present invention, DIVA can add context
9 because it understands the causes and effects in the PSE Server and how they are
10 plausibly related in a Bayesian probabilistic sense. Thus, DIVA is able to infer the
11 conditional of a hypothesized cause by reasoning backward from observed effects.
12 Indeed, DIVA can describe the prior probability of a cause, which is to say, before
13 observing any effects. As is commonly understood in the art, a prior probability is the
14 probability given only the background constraints. This is a consequence of Bayesian
15 reasoning which requires the prior probability to start the analysis.

16 The basic problem to be solved by the preferred embodiments of the present
17 invention is captured in Fig. 1A. After the PSE Server 100 completes a run, the
18 exposure profile may change significantly for any number of reasons. However, from
19 the credit analyst's point of view, the connection between cause and effect is not
20 always clear and in any case its strength cannot be accurately assessed since this
21 information is not generally available to the credit analyst.

22 According to an embodiment of the present invention, the basic idea of DIVA
23 is to correlate causes and effects, as shown in Fig. 1B, using a Bayesian network
24 which is a special knowledge base. This new approach is possible by (1) observing
25 the effect, Y_{effect} , and computing the conditional probability, $P(Y_{\text{effect}}|Z_{\text{cause}})$ and then
26 (2) assessing the plausibility of a cause, Z_{cause} , and compute $P(Z_{\text{cause}}|Y_{\text{effect}})$, provided
27 that this distribution is known through a well defined theory, empirical observations,
28 or "bootstrap" analysis. In a preferred embodiment, a combination of the latter two is

1 used, i.e., empirical observations and bootstrap analysis, to compute $P(Y_{\text{effect}}|Z_{\text{cause}})$.
2 The calculation $P(Y_{\text{effect}}|Z_{\text{cause}})$ can be “reversed” to compute $P(Z_{\text{cause}}|Y_{\text{effect}})$ using
3 Bayes theorem embodied in the Bayesian belief network.

4 Thus, according to preferred embodiments of the present invention, there is
5 provided a DIVA that uses a Bayesian belief network for systematically explaining
6 what is happening (or not happening) in the PSE Server by connecting directly
7 observable causes and effects it finds on the PSE Server. DIVA looks more deeply in
8 the data and can also look beyond the PSE Server, i.e., “over the horizon.” The
9 concept, “over the horizon,” can refer to space or time or both simultaneously.
10 In space inference, DIVA can reason about causes, for example, in the product, credit,
11 and customer information systems that are not formally part of the PSE Server but are
12 nevertheless part and parcel of the end-to-end logical flow. Accordingly, space is the
13 logical separation between independent subsystems which may or may not be
14 physically separated.

15 “Over the horizon” can be in time as well using post-diction or prediction. In
16 other words, DIVA ordinarily describes what has happened after the PSE Server
17 completes its simulation. However, it also can make predictions about what is likely to
18 happen given the incomplete information in the form of inputs from the product,
19 credit, and customer systems which must be available before the PSE Server starts its
20 simulation. This predictive feature is extremely useful because using Monte Carlo
21 simulation to measure credit risk can run for eight hours or more for just one
22 portfolio. DIVA can “forecast” the likely results before this long running process
23 starts, recommend an abort if the process looks like it won’t be successful (since the
24 inputs may look incorrect and unlikely to give accurate results), and start the next job
25 in the job stream which appears to have a greater chance of generating high quality
26 results.

27 The Bayesian belief network used by DIVA for diagnosing and/or explaining
28 changes in the PSE Server exposure profile is now described in accordance to one

embodiment of the present invention shown in Fig. 2. The Bayesian belief network 200 may be implemented by the aforementioned third-party API software. It comprises a probabilistic description of beliefs or knowledge linking causes to effects. It includes a collection of chance nodes 210 that represents the probability of change in PSE Server variables, and connections between the nodes. Table 1 defines the hypothesis variables shown in FIG. 2.

Table 1

nDeals	number of deals
nNet	number of netted deals
nPass	number of deals that could be simulated
nCef	percentage of rejected deals. Note $nCef = (nDeals - nPass) / nDeals$
dPeak	Dollar peak value used as a proxy for the exposure profile
dCmtm	dollar day-zero current mark to market
dMLIV	most likely increase in value
dCef	Credit exposure factor
xCustSys	external variable describing the Customer System (source of netting information)
xProdSys	external variable describing the Product System (source of information regarding brokered deals)
xCredSys	external variable describing Credit System (source of information for computing credit exposure factors)
_Amnts	abstract variable of high-level dollar amounts in the, e.g., day-zero CMTM
_Cnts	abstract variable of high-level counts
_Mkt	abstract variable of market data which could be observed but do not

As shown in Table 1, each node represents a random or chance variable, or uncertain quantity, which can take on two or more possible values. According to one embodiment of the present invention, the nodes represent stochastic state variables

1 that have states. In other words, the variables represent probability distributions of
2 being in a given state. In a preferred embodiment of the present invention, each node
3 has exactly two, mutually exclusive, discrete states: true or false; hence, all nodes are
4 discrete Boolean. The variables may comprise information relating to, for example,
5 input data, output data, intermediate data, and/or external data of a risk management
6 system such as the PSE Server. The arrows 220 connecting the nodes indicate the
7 existence of direct causal influences between the linked variables, and the strengths of
8 these influences are quantified by conditional probabilities. For instance, the variable
9 *dCefs* is dependent on the variable *_Amnts* in Fig. 2.

10 In a preferred embodiment of the present invention, prefixes are used in Table
11 1 to denote the type of the cause or effect being modeled. For instance, “nY” means
12 “Y is a hypothesis about counts,” and “dX” means “X is a hypothesis about dollar
13 amounts.” The other prefixed are provided in Table 2 below.

14 **Table 2**

Prefixes	Observable	Quantity
n	Yes	Count
d	Yes	Dollar
p	Yes	Proportion
v	Yes	Value
s	Yes	Structure
—	No	Abstraction
x	No	External

15
16 As shown in Table 2, there are five classes of observable variables. These
17 variables are “observable” in the sense that they can be observed and measured in the
18 PSE Server. In other words, hard evidence can be obtained for these observable
19 variables. They are the basis of “over the horizon” analysis in terms of space, time, or
20 both. In other words, the observed variables on the PSE Server can be used to infer

1 plausible causes outside the Server, as explained later in further detail. Table 2 also
2 shows two classes of unobservable variables: abstractions ($_Y$) and externals (xY). In
3 Bayesian network terminology, abstractions are called divorce variables that limit or
4 manage the fan-in of causes and effects. Fan-in herein refers to the number of parent
5 variables which affects a single variable. Abstractions serve primarily as mechanisms
6 for hiding details and organizing the network. They are devices used to help organize
7 other variables, observable or otherwise. Abstractions may also be observable
8 variables that were not chosen for observation. In this sense, abstractions are virtual
9 nodes with only circumstantial causes or effects. They are network modeling devices.
10 They cannot have hard evidence, namely, actual findings in the real world. They can
11 only have findings which are inferred from hard evidence provided elsewhere in the
12 network.

13 External variables, on the other hand, model variables in the real world except
14 they cannot be measured directly. Their existence is presumed from experience. Like
15 abstractions, external variables cannot have hard evidence, only circumstantial
16 evidence. External variables, however, are more than modeling devices. They give
17 the plausibility for systems outside the PSE Server, or in any case, outside the
18 network which is very useful information. Like abstractions, external variables only
19 have “soft” or circumstantial evidence.

20 Fig. 2 shows a Bayesian belief network 200 with only fourteen variables.
21 These variables constitute a relatively small design of low complexity chosen here for
22 simplicity in explaining the preferred embodiments of the present invention.
23 However, it should be understood that the network 200 may contain more or less
24 variables depending on the size of the PSE Servers and/or the number of variables a
25 credit analyst wishes to observe. According to an embodiment of the present
26 invention, the size and complexity of the design of the Bayesian belief network 200 is
27 a function of the number of variables in the problem domain to explain. The number
28 of nodes and their connectivity in the Bayesian belief network is a measurement of its

1 complexity. this complexity, which is called IQ , can be estimated by the following
2 formula:

$$3 \quad IQ = k - k_{min} + 1,$$

4 where k is the number of connections, and k_{min} is the minimum number of
5 connections required for a completely connected graph. For instance, the Bayesian
6 belief network of Fig. 2 has an IQ of 5.

7 Hence, the DIVA according to an embodiment of the present invention is
8 scalable to accommodate any size of the Bayesian belief network 200. The interested
9 variables in the problem domain are first order variables representing hypotheses
10 about statistically distributed causes and effects. They are used to explain a large
11 majority of exposure shifts, such as credit exposure shifts, on the PSE Server. These
12 first-order variables are chosen because they control what may be considered “first-
13 order” effects. That is, past experience indicates that when the exposure profile of the
14 PSE Server changes significantly, the expert normally considers the data from these
15 first-order variables first before looking elsewhere.

16 As mentioned earlier, connections between the nodes represent conditional
17 probabilistic influences. For example, there is a connection from a node Z
18 representing an object z to a node Y representing an object y , if Z causes Y . In such a
19 network, node Z is said to be a parent of node Y . Alternatively, node Y is said to be a
20 child of node Z . The difference between Z (big Z) and z (little z), or between Y (big
21 Y) and y (little y), will be explained later.

22 According to an embodiment of the present invention, each node and its
23 parents in the Bayesian network 200 represents a two-state conditional probability
24 distribution, namely, $P(Z_j | \mathbf{Pa}(Z_j))$, where $\mathbf{Pa}(Z_j)$ are the parent nodes of node Z_j .
25 Furthermore, the Bayesian belief network 200 represents implication, not causality.
26 Thus, if Y is a node with a parent Z , then Z *implicates* Y with probability $P(Y|Z)$. For
27 example, there is a link in the Bayesian network 200, $P(dCmtm|dPeak)$, which is
28 described as a change in the peak exposure which implicates a change in the CMTM

1 (current mark to market). In other words, if a change in peak value is observed, a
2 change in CMTM is a suspect which has to be confirmed or ruled out on the weight of
3 evidence (WOE), which will be described in detail later.

4 According to one embodiment of the present invention, the belief network 200
5 is first loaded with initial distributions or probabilities consistent with the state of
6 knowledge prior to considering evidence. In other words, the belief network 200 is
7 initially biased in favor of certain conclusions. The source of this initial bias may
8 range from an objective, well-defined theory to completely subjective assessments.

9 According to an embodiment of the present invention, the initial distributions
10 of variables x and y are hypotheses, as denoted by $H(x)$ and $H(y)$, respectively. Then
11 a node x with a parent y specifies a hypothesis $H(x)$ given $H(y)$ written as $H(x)|H(y)$.
12 $H(x)$ is the working or null hypothesis about x , namely, that “ x has not changed.”
13 Thus, the initial distributions have been set up such that the bias is toward disbelief
14 about changes which in fact corresponds to direct experience because, as noted
15 earlier, most variables in the PSE Server are correct most of the time. Thus, the null
16 hypothesis has a practical basis in reality. As is understood in the art, a null
17 hypothesis is one that specifies a particular state for the parameter being studied. This
18 hypothesis usually represents the standard operating procedure of a system of known
19 specifications.

20 Hypotheses, of course, are statements. They are either true (T) or false (F),
21 and they obey the rules of logic. Because $H(x)$ is the working hypothesis, it is
22 initially assumed to be true. Thus, for the sake of simplicity, $H(x)$ herein means
23 $H(x)=T$. Then $\sim H(x)$ negates the assumption, meaning the hypothesis that “ x has not
24 changed” is false. $H(x)H(y)$ means the hypothesis that “ x has not changed” and “ y
25 has not changed” is true. $H(x) + H(y)$ means the hypothesis that “ x has not changed”
26 or the hypothesis “ y has not changed” is true or both are true.

27 Because the hypotheses are logical, the nodes 210 in the belief network 200
28 shown in Fig. 2 are two-state or Boolean, as mentioned earlier. That is, each variable

has only two possible states: T or F. The Bayesian belief network is now used to determine the probability of the null hypothesis for each variable. In classical statistics, this the meaning of the p-value: the probability of incorrectly rejecting the null hypothesis. Consequently, the p-value of $H(x)$ can be written as $P(H(x))$.

When the null is conditioned, for example, then the conditional working hypothesis about x is true given that some other hypothesis about y is true. As mentioned earlier, this is denoted by $H(x)|H(y)$. Consequently, the conditional probability is $P(H(x)|H(y))$, that is, the probability of the hypothesis that “ x has not changed” given the hypothesis that “ y has not changed”. To avoid confusion with the notation and without loss of generality, $P(X|Y)$ will be used hereinafter to denote the conditional probability, wherein it is understood that X and Y are hypotheses about x and y , respectively. In other words,

$$P(X|Y) = P(H(x)|H(y)); \text{ with}$$

$$X = H(x),$$

$$Y = H(y).$$

It should be clarified that X and Y are not random variables in the classical sense. What is distributed is not X or Y but the probability $P(X|Y)$. Hypotheses X and Y are logical statements about objects x and y , and $P(X|Y)$ is a plausible statement about the believability of X assuming Y .

According to an embodiment of the present invention, the design of the Bayesian network comprises two features: quality and quantity. Quality is expressed in the structure or architecture of the network while quantity is expressed by the probability distributions. The quality or network structure is the more important feature of the two, for it describes the precise nature of believed implications in the system. Thus, $P(X|Y)$ gives a different implication relationship compared to $P(Y|X)$.

For instance, referring back to Fig. 2, let $dCmtm$ 212 represent the hypothesis that the “current mark to market exposure of the portfolio has not changed,” and let $dPeak$ 214 represent the hypothesis that the “dollar peak exposure value of the

1 portfolio has not changed.” Thus, $P(dCmtm|dPeak)$ and $P(dPeak|dCmtm)$ are
 2 permissible by the rules of logic, but in practice they have different meanings. The
 3 former is meaningful for implication as a weak form of causality and is used in
 4 preferred embodiments of the present invention. The latter is meaningful for a strong
 5 form of causality which is not advocated because while $dCmtm$ 212 $dCmtm$ 212 does
 6 effect $dPeak$ 214, the nature of this relationship is unreliable for purposes of the
 7 present invention.

8 Another reason that the network structure is more important is that given
 9 sufficient evidence, a Bayesian network can converge to the “right” answer despite its
 10 initial bias. “Right” in this case is used in the sense of “same.” Convergence and the
 11 rate of convergence depends on the network’s initial bias as well as on the WOE that
 12 has been submitted. Theoretically, this is proven by the observation that the initial
 13 bias acts as a constant or level and in the limit the ratio of the two systems of beliefs
 14 equals one because the WOE’s are the same, overriding the initial discrepancy. The
 15 mathematical justification for this goes as follows.

16 Let $O(A_{ik})$ be the prior odds of some hypothesis A_i under a belief system k .
 17 Let $O(A_{ij})$ be the prior odds for the same hypothesis A_i under a belief system j .
 18 Systems k and j differ only in the prior probabilities; however, they agree on the
 19 meaning of evidence given in the Bayes factor, β_i . Thus, given sufficiently large
 20 evidence, the WOE for the two systems will converge, i.e.,

$$\lim_{\beta_i \rightarrow \infty} \frac{\log O(A_{ik}) + \beta_i}{\log O(A_{ij}) + \beta_i} = 1$$

22 Thus, while the choice for the initial distributions is not of primary concern, such
 23 distributions should be chosen carefully to avoid distributions that cause the belief
 24 network to contradict itself.

25 Self-contradiction by the belief network may ultimately cause problems. This
 26 is an issue that involves Gödel’s incompleteness theorem, as mentioned earlier. The
 27 solution is Cromwell’s Rule, which forbids the use of zero or one probabilities

anywhere in the Bayesian network, including initial probabilities. Cromwell's Rule also plays a special role when re-sampling is used to generate the likelihood distribution, $P(f_i|A)$. This will be discussed later.

According to an embodiment of the present invention, the initial distributions or probabilities comprise prior probabilities and initial conditional probabilities. The initial probabilities can be set by (a) using the advice of an "expert," (b) learning from the data automatically, or (c) applying the following values (which may be justified by observing again that most of the data is correct most of the time):

$$P(Z_j=T|Z_{k \neq j}=T)=0.95;$$

$$P(Z_j=T|Z_{k \neq j}=F)=0.05$$

The first distribution indicates a 95% certainty that the null hypothesis is correct, i.e., the feature represented by Z_j has not changed when its parent, $Z_{k \neq j}$, has not changed.

The second distribution indicates a 5% certainty that the null hypothesis is correct, i.e., the feature represented by Z_j has not changed when its parent, $Z_{k \neq j}$, has changed.

This follows from common sense and conforms, once again, to actual experience.

When Z_j has more than one parent, then the initial conditional probabilities can be derived from noisy-or functions or logical-or functions. If, for instance, a network $P(A|B,C)$ is built using noisy-or, the CPT can be calculated using:

$$P(A | B C) = P(A | B) + P(A | C) - P(A | B) P(A | C),$$

where $A=T$ represents some probability conditioned on $B=T$ and $C=T$. In other words, each hypothesis is in the true state. When a hypothesis is not in the true state, namely, $A=T$, $B=T$, and $C=F$, the CPT is calculated using:

$$P(A | B C) = P(A | B),$$

$$P(A | B) = P(B);$$

and when $A=T$, $B=F$, and $C=T$, the CPT is calculated using:

$$P(A | B C) = P(A | C),$$

$$P(A | C) = P(C);$$

and when $A=T$, $B=F$, and $C=F$, the CPT is calculated using:

$$P(A | B C) = 1 - [P(A | B) + P(A | C) - P(A | B) P(A | C)],$$

According to an embodiment of the present invention, the noisy-or calculations are used for two important reasons. First, the noisy-or can be generalized for an arbitrary number of parents where conditional probabilities can be combined using set theoretic permutations. Thus, for $P(A|BCD)$, the probabilities may be combined as

$$P(A|BCD) = P(A|B) + P(A|C) + P(A|D) - [P(A|B)P(A|C) + P(A|B)P(A|D) + P(A|C)P(A|D)] + P(A|B)P(A|C)P(A|D),$$

for the case where all hypotheses are in the true state.

Second, noisy-or satisfies Cromwell's Rule because the resulting probability will be asymptotically one (i.e., $\sum P(A|Pa(A)) \rightarrow 1$) as long as the conditional probabilities are not zero or one where $Pa(A)$ are the individual parents of A. If the network $P(A|BC)$ is built using logical-or, there is no need to calculate the above conditional equations. In fact, logical-or networks are much simpler to construct. However, they do not satisfy Cromwell's Rule because by definition the CPT will contain a zero probability if all hypotheses are in the false state. The network will contain a one probability otherwise. This need not be a problem. As long as the prior probabilities are Cromwellian (i.e., non-zero and non-one), contradictions can be avoided.

To make the distinction between noisy-or and logical-or clear, illustrative CPTs for both noisy-or and logical-or are given in Tables 3 and 4 below for a network example, $P(A|BC)$. In either case, the prior probabilities are set at, for example, $P(B=T)=0.85$ and $P(C=T)=0.95$. Note: $P(B=F)=1-P(B=T)=0.15$ and $P(C=F)=1-P(C=T)=0.05$. First, the values for noisy-or CPT are calculated using the above equations as:

Table 5

BC		FF	FT	TF	TT
A BC	F	0.9925	0.05	0.15	0.0075

	T	0.0075	0.95	0.85	0.9925
--	---	--------	------	------	--------

1

2 As shown from Table 5, the initial conditional probabilities are determined from the
 3 prior probabilities. However, the identical configuration under logical-or is:

4

Table 6

BC		FF	FT	TF	TT
A BC	F	1	0	0	1
	T	0	1	1	1

5

6 Thus, logical-or and noisy-or are not identical. However, as the two CPTs
 7 suggest above, they can serve as approximations for each other. In general, noisy-or
 8 is preferred when the fan-in is low, and logical-or is preferred when the fan-in is high.
 9 When fan-in is low, the above equation can be readily calculated and verified. When
 10 the fan-in is high, the above equation can be calculated but the number of
 11 combinations is high. Moreover, even if the calculation is automated, it will remain
 12 difficult to verify each combination of inputs. For instance, for a node with eight
 13 parents, there are 2^N or $2^8=256$ combinations (because each node has two states).
 14 Also, because the noisy-or probabilities still must be entered manually into a causal
 15 probability table (CPT), changing the probability of one of the parents, i.e., B in
 16 $P(A|B)$, will affect the entire network. This is impractical if the fan-in is highly.

17 A DIVA that uses the aforementioned Bayesian belief network for analyzing
 18 the PSE Server is now described. Fig. 3 shows a DIVA architecture 300 according to
 19 an embodiment of the present invention. The DIVA 300 comprises programs, data,
 20 and a knowledge base. The programs are written in two modules, a normative auto
 21 assistant (NAA) 310 and a data grabber (not shown). The term “normative” herein
 22 refers to the reliance on underlying mathematical theories, such as the laws of
 23 probability. The NAA 310 is where all the Bayesian logic is programmed. It can be

1 implemented by any suitable computer programming language, such as Microsoft
2 Visual C++. Thus, the NAA 310 can run wherever there is a compiler for the
3 computer programming language. The data grabber gets the raw data of the
4 observable variables in the PSE Server for the NAA 310. According to an
5 embodiment of the present invention, the data grabber can be written in a program
6 script, such as Perl, and runs on the PSE Server.

7 According to a further embodiment of the present invention, the two major
8 components of the NAA 310 are the electronic brain equivalent (EBE) 312 and the
9 main evidence extraction component (MEECO) 314. Each of these are programming
10 objects, such as C++ objects, that interact with each other in a tight loop as shown in
11 Fig. 3. The main function of the EBE 312 is to thinly encapsulate using object-
12 orientation calls to the API of the third-party API software, which is not object-
13 oriented. The EBE 312 further provides mapping between three name spaces: nodes,
14 variables, and observables.

15 Nodes are objects which the API manipulates as opaque types. The API
16 software also has *domains*, objects that describe a Bayesian network which contains
17 nodes. The EBE 312 completely hides these details. Variables are objects of interest,
18 that is, the fourteen variables given in the tables above. Observables are a subset of
19 variables, i.e., those given in the table of observable variables. The distinction
20 between one another name space is needed for two reasons.

21 First, variables are a construct invented as a proxy for the Bayesian network
22 nodes. These nodes are C pointers in the third-party API software, whereas variables
23 are integers. Indeed a variable is just an index to a vector of void pointers. Second,
24 the ordering of the variables is arbitrary: the Bayesian network nodes are organized
25 abstractly (i.e., the algorithm of assignment is hidden in the API software) and as the
26 nodes are loaded, they are assigned an integer index in a sequence. Thus, mapping is
27 needed between variables and nodes.

1 Second, as a consequence, observables are scattered among the variables in
2 random sequence, although observables are generally manipulated in a given order
3 according to a speculative hypothesizer or interpreter (ASH) function that may be
4 implemented implicitly by the NAA 310. This ASH function will be discussed later.
5 Thus a mapping is needed between variables and observables. The EBE 312 manages
6 this. The relationships between these name spaces are shown in Fig. 4.

7 As mentioned earlier, the MEECO 314 is also a programming object. Its
8 primary function is to convert raw data of the observable variables into evidence.
9 Implicitly encapsulating a weigh-in (WEIN) function, the MEECO 314 then sends the
10 evidentiary findings into the EBE 312. This WEIN function will be discussed later.
11 The EBE 312 also retrieves beliefs by variable from the Bayesian belief network 320
12 whether or not "hard" evidence has been entered. If no evidence has been supplied,
13 the EBE 312 returns the initial priors and conditionals. As also shown in Fig. 3, the
14 NAA 310 interacts with a fast recursive diagnostic (FRED) interpreter 360, via a
15 confirmation matrix 350. The FRED interpreter 360 may be a separate program, as
16 shown in Fig. 3, or it may be an object embedded within the NAA 310. The
17 algorithm for FRED interpreter 360 is provided and discussed next in accordance to
18 an embodiment of the present invention.

19 The FRED algorithm automates the interpretation of the confirmation matrix.
20 It can be easily programmed and used to write a more systematic report for the user.
21 The idea of FRED testing the "complexity" of the matrix and analyzing the
22 confirmations accordingly.

23 The complexity, K , is an estimate of the interpretation effort. It is the number
24 of self-confirmations ≥ 5 db, not including the peak exposure.

25 FRED works recursively using K . At any given level of recursion, FRED
26 wants to interpret matrices of low or moderate complexity. If the complexity is
27 greater, it reduces the complexity by one and calls itself recursively, trying again. It
28 then backtracks.

1 The FRED algorithm is given below. On the notation, $[V]$ is a vector of
 2 variables, $n([V])$ is the length of the vector, and $[V]$ starts at index 0. $V_i \rightarrow V_j$ means
 3 variable i implicates variable j or alternatively, variable j effects variable i .

```

4  procedure fred([V])
5  begin
6    K = n([V])
7    case K ≤ 1:      // low complexity
8      report the  $V_0$  as the explanation with confirmation
9      check unobservables and report indirect confirmations ≥
10 5 db
11    return
12    case 1 < K ≤ 2: // moderate complexity
13      sort [V] by implication using the BN
14      if  $V_1 \rightarrow V_0$  then
15        fred([ $V_0$ ])
16      else if  $V_0 \rightarrow V_1$  then
17        fred([ $V_1$ ])
18      else // two possible effects, neither implicating the
19 other
20        Sort [V] by marginal importance
21        fred([ $V_0$ ])
22        fred([ $V_1$ ])
23    case K > 2:      // high complexity
24      Sort [V] by implication using the BN
25      if  $V_j \rightarrow V_i$  for all  $i \neq j$  then
26        fred([ $V_j$ ])
27      else // there are two or more effects
28        Sort [V] by self-confirmation
29        fred([ $V_0 \dots V_{n-2}$ ]) // eliminate the lowest confirmation
30        fred([ $V_{n-1}$ ])       // backtrack to explain eliminated
31 variable
32 end procedure fred
33

```

34 Note that the FRED algorithm does not take into account potential
 35 inconsistencies. For instance, there's positive self-confirmation for *dCef* but no self-
 36 confirmation for *dCmtm* nor for *dMliv*. Technically this is a data conflict which
 37 should be written into the algorithm.

38 According to an embodiment of the present invention, the raw data of each
 39 observable variable comprise two types: bias data 330 and fact data 340. Bias data

1 are historical views of what has happened in the past which bias the analysis. The
2 fact data are the data to be explained. The biases 330 and facts 340 comprise $k_r \times N$
3 tables of raw data extracted from the PSE Server via a server archive (not shown),
4 where N is the number of observable variables which is 8 for the Bayesian belief
5 network 200 of Fig. 2. (Actually, the raw data contains N=7 variables but N=8 are
6 created by deriving one of the variables, nCef, from two others.) The value of k_r , i.e.,
7 the number of rows or vectors of variables, is independent for the biases and facts.

8 The knowledge base of DIVA comprises the Bayesian network 200 (Fig. 2) as
9 implemented by the aforementioned third-party API software. Thus, the knowledge
10 base includes all observable and unobservable variables, the network of conditional
11 probabilities, and the initial priors and conditional parameters.

12 Fig. 3 is a specific embodiment of Fig. 5. In other words, Fig. 5 shows a more
13 general scheme for a DIVA architecture in accordance with preferred embodiments of
14 the present invention. Fig. 5 depicts a general DIVA architecture 500 showing the
15 main functional modules and their relationships in accordance to another embodiment
16 of the present invention. These modules represent a plurality of support features
17 which DIVA may contain to effectively use the Bayesian belief network as
18 implemented by the API software.

19 As shown in Fig. 5, the belief network is loaded and accessed through the
20 belief network API of the API software using an EBE 520 of DIVA. The EBE 520 is
21 the same EBE 312 shown previously in Fig. 3. The EBE 520 also takes as input the
22 evidence from the weigh-in (WEIN) 510, gives its data to the Bayesian belief network
23 (not shown) to update the state of knowledge, and gets back beliefs which it then
24 sends to an Automated Speculative Hypothesizer (ASH) 560 to interpret. The
25 Bayesian belief network used for the DIVA 500 is the same network used in the
26 DIVA 300 of Fig. 3. The ASH 560 then sends the prospects according to its
27 interpretation of the beliefs to the Main Evidence Extraction Component (MEECO)

1 530. The relationships between the WEIN 510, the ASH 560, and the MEECO 530
2 are described next.

3 As mentioned earlier, the automated speculative hypothesizer or ASH 560
4 interprets beliefs from the EBE 520. In other words, the ASH 560 determine the new
5 evidence to extract from the PSE Server. The ASH 520 may be a programming
6 object used for applying the constraints 550 for seeking out the most plausible suspect
7 which has not already been implicated or ruled out. The issue to be considered is the
8 classic one of searching depth-first vs. breath-first. In other words, according to one
9 embodiment of the present invention, the ASH 560 can output the top N prospects of
10 interpreted beliefs and let the DIVA system try to absorb them all in one evidence
11 instantiation. Alternatively, the ASH 560 can output one prospect at a time to allow
12 the DIVA system to absorb each in turn before a new prospect is considered. The
13 DIVA system can advance along a specific path, eliminating variables in a pre-
14 programmed manner. This is called structured supervision. Alternatively, the DIVA
15 system can jump to conclusions given whatever it finds interesting. This is called
16 unstructured supervision.

17 As mentioned earlier, the above options and others are decided by constraints
18 550. In a preferred embodiment, the Jaynes' sequential admission rule is applied as a
19 constraint. This rule provides for the testing of the most promising prospect(s) first
20 and then proceeding to the next promising one(s). Thus, this implies that the ASH
21 560 may sort all beliefs into ascending order and pick the top one(s) to pursue.

22 Referring back to the DIVA architecture 300 of Fig. 3. Although there is not
23 shown an ASH or speculative interpreter in the loop between the EBE 312 and the
24 MEECO 314, the aforementioned ASH function remains in the NAA 310 in
25 accordance to that embodiment of the present invention. Specifically, the plausibility
26 constraint (as depicted by constraints 550 in Fig. 5) can be removed, and the NAA
27 310 can be programmed to seek out suspects in a pre-programmed manner.
28 According to DIVA architecture 300 of Fig. 3, the NAA 310 is sufficiently fast such

1 that all variables can be checked without serious time penalties. Thus, it is redundant
2 to use an ASH to optimize the search by going after the most promising prospects in
3 the DIVA 300.

4 Reference is now made to the Main Evidence Extraction Component or
5 MEECO 530 in Fig. 3. As seen from the figure, the MEECO 530 takes the prospects
6 output by the ASH 560 and by searching the PSE Server archive 540 for raw biases
7 and fact data of observable variables, converts the prospects to factoids. A factoid
8 includes factual data of an evidentiary nature that remains to be substantiated.

9 The MEECO 530 extracts factoids by analyzing changes in the PSE Server
10 historical backup. If the MEECO 530 is given a list of backups, it produces a baseline
11 statistical database, which contains the sum of squares for each variable. If it is given
12 just two backups, it produces just the changes between two runs. According to a
13 preferred embodiment of the present invention, the MEECO 530 extracts everything;
14 however, it does not use thresholds. That is the job for the WEIN 510. It should be
15 noted that the MEECO 314 of the DIVA architecture 300 (Fig. 3) is similar to the
16 MEECO 530 of the DIVA architecture 500, except that the MEECO 314 also
17 performs the job of the WEIN 510, which is described next.

18 The WEIN 510 is a crucial component of DIVA. It allows DIVA to find the
19 needle in the haystack as follows. DIVA keeps sufficient statistics in a database
20 which is built and updated periodically by the MEECO 530. To diagnose a feed,
21 DIVA invokes the MEECO 530 for the prior and current run and extracts the one-run
22 factoids. The WEIN 510 then weighs these factoids using statistical re-sampling and
23 calculates the conditional for the given factoid. This conditional is the probability of
24 the null hypothesis, namely, of obtaining the given factoid assuming it does not
25 represent a significant change. The conditional for a given factoid f_i for a variable i ,
26 as denoted by a node in the Bayesian belief network 200 (Fig. 2) is mathematically
27 denoted by:

$$P(f_i | A_i)$$

1 where A_i is a working hypothesis for the variable i .

2 The distribution, $P(f_i|A)$, must be treated carefully when re-sampling. The
3 main issue is simply that f_i may not exist in the distribution because re-sampling
4 creates only a range of elements. In particular, f_i may exceed the last element in the
5 re-sampled distribution or it may precede the first element in the distribution. It
6 would be simple to set the probabilities to one and zero respectively but then would
7 not satisfy Cromwell's Rule. Thus, when f_i is larger than the last element, v_N , then

$$8 \quad P(f_i|A) = 1/[N (1+(f_i - v_N)/v_N)]$$

9 When f_i is smaller than the first element, v_0 , then

$$10 \quad P(f_i|A) = 1-1/[N (1+(v_0 - f_i)/v_0)]$$

11 N is the size of the re-sampled distribution.

12 The WOE, i.e., the evidence obtained by the WEIN 510 weighing the factoids
13 is then given by the Bayes factor,

$$14 \quad \beta_i = \log \frac{P(f_i | A_i)}{P(f_i | \sim A_i)}$$

15 which is the log of the likelihood ratio. DIVA does not have direct access to $P(f_i | \sim A)$
16 because generally the credit analyst rejects all $\sim A$ data feeds. Therefore, $P(f_i | \sim A)$
17 may be estimated as follows. It is conventionally known in the art that credit analysts
18 tend to reject f_i when it seems obviously less than a threshold value v , which is chosen
19 in accordance to business rules. This estimation can be simulated by computing the
20 transformation,

$$21 \quad P(f_i | \sim A) \approx P(g_A(f_i) | A)$$

22 where g is the rescale functional. The rescale functional can be any function.

23 However, for the sake of demonstration and simplicity, g is chosen such that

$$24 \quad K_A f = g_A(f)$$

25 where K_A is the rescale factor which depends on A . In this case, the factoid is scaled
26 linearly; however, the probability distribution, $P(f_i | A)$, is non-linearly transformed.

27 K_A is chosen in such a way that it stretches $P(f_i | A)$ and the resulting β_i approximately

1 follows the credit analysts business rules. Business rules describe when and under
 2 what conditions f_i should be rejected. Typically, f_i is rejected when it exceeds the
 3 business threshold, namely, v .

4 Factoids need to be rescaled because, again, the $P(f_i | \sim A)$ distribution is not
 5 available but which is needed for the WOE calculation. Thus, $P(f_i | \sim A)$ may
 6 estimated using the rescale technique.

7 According to an embodiment of the present invention, the above calculations
 8 for the Bayes factor β_i are done using the Monte Carlo simulation as implemented by
 9 the MEECO 314 shown in Fig. 3, or alternatively, by the WEIN 510 shown in Fig. 5.

10 The third-party API software does not use β_i directly. Instead, it uses the
 11 likelihood ratio of β_i to calculate the posterior probability $P(A_i | f_i)$ using the odds form
 12 of Bayes' Rule, namely,

$$13 \quad O(A_i | f_i) = O(A_i) \frac{P(f_i | A_i)}{P(f_i | \sim A_i)}$$

14 wherein,

$$15 \quad O(A_i) = P(A_i) / P(\sim A_i), \text{ and}$$

$$16 \quad O(A_i | f_i) = P(A_i | f_i) / P(\sim A_i | f_i)$$

17 and presents to the credit analyst the *confirmation* which is measured in decibels,
 18 namely,

$$19 \quad C_{i,j} = 10 \log_{10} \frac{P(A_i | f_j)}{P(A_i)}$$

20 which is just the ratio of the posterior probability to the prior probability; wherein, A_i
 21 is the working hypothesis for variable i , and f_j is the factoid for variable j .

22 As explained earlier, the above confirmation equation is derived from the
 23 Bayes factor. In other words, when a finding is entered into the belief network, the
 24 API software propagates the evidence to all nodes. Recall an earlier discussion that
 25 the API software uses special mathematical methods and system techniques to make
 26 this feasible because the complexity $O(2^N)$ time is otherwise unreasonable. DIVA has

prior probabilities from the initial priors and conditionals. It receives the posterior probabilities $P(A_i|f_j)$ from the updated beliefs, which the EBE 312 generates. Thus, DIVA can compute the confirmation.

The above equation shows that $C_{i,j}$ is the log change in probability of a variable in response to evidence about another variable. Thus,

- If $C_{i,j} > 0$, the working hypothesis, A_i , is supported by the evidence. In other words, A_i is confirmed.
- If $C_{i,j} < 0$, then A_i is denied by the evidence. It is disconfirmed.
- If $C_{i,j} = 0$, then A_i is neither supported nor denied by the evidence.

According to an embodiment of the present invention, there is concern with only the second case where $C_{i,j} > 0$, and only when $C_{i,j} \geq 5$ because this is the threshold of “positive” confirmation of A_i . Above about 11 decibels there is “strong” confirmation of A_i , and above about 22 there is “decisive” confirmation of A_i . Table 5 shows the commonly known scientific standards of evidence as developed by the British geophysicist, Sir Harold Jefferys, back in the 1930s, as applied in an embodiment of the present invention.

Table 5

Confirmation (db)	Evidence for A_i
< 0	None; evidence against A_i
$= 0$	Inconclusive
$> 0-5$	bare evidence
5-11	Positive
11-22	Strong
> 22	Decisive

Referring back to Fig. 3, the NAA 310 of DIVA computes a *confirmation matrix* 350 from the above confirmation equation. This matrix is the main

1 interpretive report used to “explain” the exposure shifts. According to an
 2 embodiment of the present invention, programmable rules are then provided in DIVA
 3 to interpret the matrix 350. Moreover, the matrix 350 is numerical.

4 The matrix 350 provides hard confirmation along the diagonal and
 5 circumstantial confirmation off the diagonal. In other words, C_{ii} , is the hard
 6 confirmation for finding i on observed variable i . This is also called self-
 7 confirmation. The circumstantial confirmation, C_{ij} , gives the “soft” effect of finding i
 8 on variable j which may be observable or unobservable. This is also called cross-
 9 confirmation. Because there are observables and unobservables in the Bayesian belief
 10 network 200 (Fig. 2), the matrix 350 includes two sub-matrices. The top sub-matrix
 11 comprises a $k \times k$ square matrix, and includes the observable variables. This top sub-
 12 matrix indicates how much the self-evidence confirms or denies the working
 13 hypothesis, namely, that some variable A_i has not changed. As mentioned earlier, a
 14 meaningful positive value (≥ 5) along this diagonal indicates the data is suggesting a
 15 significant change in the corresponding observable variable.

16 With regard to the off-diagonal values in the top sub-matrix, these are
 17 indications of sensitivities change logically prior to considering the self-evidence for
 18 the respective variable. In other words, C_{ij} for $i \neq j$ confirms (or denies) the potential
 19 impact of evidence for variable A_j on variable A_i . The impact is potential because
 20 until the evidence on A_i is actually reviewed, there is only indirect confirmation as
 21 opposed to direct confirmation. As for the bottom sub-matrix, it comprises a $m \times k$
 22 rectangular matrix for m unobservable variables. These elements are all off-diagonal
 23 and thus the confirmations are all circumstantial.

24 While looking at individual entries in the confirmation matrix is definitive, it is
 25 sometimes helpful to see the big picture of implications in a risk management system
 26 such as the PSE Server. For this, the concept of *importance* is used of which there are
 27 several varieties. Table 6 shows the importance measurements in accordance to an
 28 embodiment of the present invention.

Table 2

Importance	Measurement
Self-importance	$\gamma_j = c_{j,j}$
Marginal importance	$\gamma_j = \sum_i^k c_{i,j}$
Absolute marginal importance	$\gamma_j = \sum_i^{k+m} c_{i,j}$
Relative importance	$\gamma_{j,l} = \gamma_j - \gamma_l$

According to an embodiment of the present invention, a generic mode of DIVA operation is essentially assumed. There are, however, specific constraints or "factory settings," that can tailor DIVA for particular operative environments. These setting are shown in Table 7 below.

The primary differences between the settings involve the initiation and the confirmation *credibility threshold*. In the "real-time" setting, DIVA is automatically invoked by a decision check on the hold/release cycle. In the "follow up" and "passive excesses" settings, the credit analyst invokes DIVA manually. Finally, on the "deep six" setting, DIVA is run periodically to "scrub" the system's data feed.

The credibility threshold is the credibility level below which DIVA suppresses explanations of the confirmation matrix. The point is to qualify or filter explanations in a way that is consistent with the operative environment. For instance, in the real-time mode the credit analyst must in a timely manner decide whether to hold or release a feed. The quality of an explanation, namely its credibility, should be consistent with the criticality of the situation. Thus, DIVA reports only the strongest explanations during real-time.

Table 7

Setting	Mode	Explanation Objective	Initiated	Credibility threshold
Real-time	On-line	Changes in exposure profile during hold/release phase	Decision check	Strong
Follow-up	Off-line	Changes in exposure profile following up the hold/release phase	On demand	Strong
Passive excesses	Off-line	Persistent features in exposure profile	On demand	Substantial
Deep Six	Off-line	Potential problems buried deep in the data	Cron (UNIX utility)	Bare mention

DIVA uses a normative, rather than descriptive, approach to explaining the PSE server. It models how the system behaves and not how the credit analyst behaves. Thus DIVA is a tool for logical analysis. It is designed to support, rather than replace, the credit analyst.

Although only a few exemplary embodiments of this invention have been described in detail above, those skilled in the art will readily appreciate that many modifications are possible in the exemplary embodiments without materially departing from the novel teachings and advantages of this invention. Accordingly, all such modifications are intended to be included within the scope of this invention as defined in the following claims. Furthermore, any means-plus-function clauses in the claims (invoked only if expressly recited) are intended to cover the structures described herein as performing the recited function and all equivalents thereto, including, but not limited to, structural equivalents, equivalent structures, and other equivalents.